09/13/02

A DISCUSSION OF THE QUALITY OF ESTIMATES FROM THE AMERICAN COMMUNITY SURVEY FOR SMALL POPULATION GROUPS

Charles H. Alexander

U.S. Census Bureau

August 26, 2002

This version is a draft prepared for discussion at the Fall 2002 Census Advisory Committees Meetings

EXECUTIVE SUMMARY

Census stakeholders have raised questions about the quality of estimates for small population groups from the American Community Survey (ACS). Some are concerned about how to use the annually updated multi-year averages the survey products will show for populations of less than 65,000. This is especially the case for the 5-year averages the Census Bureau proposes to use for smaller groups to replace the decennial long-form "snapshot." This report focuses on these and other concerns about how the survey design and operations affect the quality of the statistics for smaller population groups. We explain how we are addressing the issues. It explains why the Census Bureau and many data users view the annually updated estimates from the ACS to be particularly advantageous in providing clearer information about small groups.

The major questions and the responses are:

Question 1: *What is the impact of having a smaller ACS sample size in any single 5-year period than the long form has in the census year?*

Response: The increase in the margin of error in the ACS, compared with the long form, is modest and the report cites examples. There are potential improvements in reduced nonsampling errors that is expected to compensate for the larger margins of error. This report demonstrates, with examples of change from 1990 to 2000, the advantage the ACS provides small population groups by being able to track the direction, timing, and level of what are oft-times large changes in the population and characteristics of small racial, Hispanic origin, and ancestry groups. We conclude that overall, the ACS design benefits small population groups by improving the richness of the information available to them.

Question 2. *Are the multi-year averages more difficult to interpret than point-in-time estimates?*

Response: If the characteristics of a population change substantially during the time period covered by the average, getting some information every year is better than getting information for one moment in time every 12 years. If there is little change in the population over the time covered by the average, the interpretation is about the same as that of a point-in-time estimate with the advantage that the ACS estimate is more current than the historical decennial census long-form estimate. The report gives examples and illustrates how to analyze 1-, 3-, and 5-year to learn the nature and timing of changes.

Question 3. *Won't there be an increase in the standard errors in areas where response by mail or telephone is relatively low because you use a subsample of 1 in 3 nonresponse cases for follow up with personal visits from Field Representatives to collect the data?*

Response: It is important to bring the standard errors for all groups and areas in line with the objectives for response rates overall. The Census Bureau proposes to address this issue through several techniques, including (a) using a subsampling rate larger than 1 in 3 in areas with low mail response; and (b) making it easier for people with limited English proficiency to respond by mail or telephone. Nonsampling errors can be larger than sampling errors and so our research program monitors both. Because the Field

Representatives are experienced, they have had good success in areas where it has traditionally been difficult to collect survey statistics during the short decennial census operations with temporary staff. This may partially compensate for subsampling the nonresponse cases for follow up to collect responses to the questionnaire.

Question 4. *How can a small sample be representative of a small population group that is geographically dispersed?*

Response: The report describes how statisticians use a statistical concept, the “ law of averages,” to calculate the “margin of error” or “confidence interval” for the ACS estimates. While the estimates for a single month may be very unpredictable, data averaged over 60 months provides reasonably stable estimates.

I. INTRODUCTION

This report discusses the quality and usefulness of estimates from the planned American Community Survey (ACS), for very small populations groups. The ACS is intended to replace the long form survey in the 2010 census. Since the long form is unique as a source of information about smaller population groups, a priority objective of the ACS design has been to provide good information about smaller groups.

The general premise of the ACS design is that by spreading the "long form" sample across the decade, it is possible to provide updated information for all sizes of population groups. In principle, this should be especially advantageous for small population groups, because there is currently very little information about how these populations change over time. Also, the ACS is expected to have more consistent quality because of the advantages of a continuous operation, which is especially important for those small groups that have traditionally been difficult to include in surveys and collect information about characteristics.

Questions have been asked about the quality of ACS estimates for very small population groups. These concerns are described in Section III, with responses in subsequent sections. We have described the ACS as replacing the long form "snapshot" with a "video." Using this metaphor, the most widespread concerns are

(1) that a “freeze frame” from the video is not as clear as a snapshot, and (2) that if the subject of the picture is small and fast-moving, the video may show a blur.

The response is, continuing the metaphor, that the freeze frame is almost as clear as the snapshot, and provides the advantage of being able to look at a freeze frame at any time. For fast-moving subjects, a video at least tells you that the subject is moving and in what direction, while a snapshot misses the action totally. Small population groups have the potential to change more dramatically than larger groups, so having a “video” is particularly valuable for smaller groups.

The sections and appendices of this paper present the basic issues, as well as some of the more complex statistical issues.

Comments on topics for which this report may not have effectively explained the issues are welcome. There may be subsequent revisions of the report based on discussions with the Census Bureau Advisory Committees, or as new information is available from evaluations of the ACS and comparisons with Census 2000.

II. BACKGROUND ON THE OPERATIONS AND DESIGN OF THE AMERICAN COMMUNITY SURVEY

The ACS is part of a plan to re-engineer the 2010 census. Besides replacing the long form with the ACS, the plan includes modernizing the geographic system and updating the list of addresses throughout the decade (i.e., MAF/TIGER), and early planning and research to design better and more accurate ways to count the population in 2010.

The ACS plan is to start in 2003, with an annual sample of 3 million addresses spread across the list of addresses in each census tract, covering all places (such as cities or towns), American Indian Reservations, Alaska Native villages, and Hawaiian Homelands. About 250,000 addresses will be contacted for the first time each month. No address will be in sample more than once in a 5-year period. We expect that for most addresses, there will be about forty years between ACS interviews.

Most addresses in the sample start out with a mail questionnaire in their first month, with a prenotice, a reminder card, and a targeted second mailing. In the second month, Census staff follow up at addresses that did not respond and for which a telephone number is available with a Computer-Assisted Telephone Interviewing (CATI) operation. In the third month, we select a one-in-three sample of addresses which have still not responded for follow-up by Field Representatives who use Computer-Assisted Personal Interviewing (CAPI). Mail responses with substantial amounts of missing data are designated for recontact by telephone in a “failed edit follow-up” operation. Units for which there is no usable mailing address skip the mail and CATI phases. A two-in-three sample of such units goes straight to the CAPI operation.

As was done for the last three census long form samples, small governmental units will be sampled at a higher rate, depending on the population of the area. In particular, the smallest governmental units will be sampled at a rate of 10 percent per year. Addresses in large census tracts are sampled at a somewhat lower rate, unless they are in a small governmental unit.

The survey designers are considering a plan to oversample census tracts that have much lower-than-average mail response rates. If this plan is implemented in such areas, the CAPI follow-up rate would be greater than 1 in 3. To make up for this, the initial sampling rate would be reduced slightly in tracts with above-average mail response rates. Options for doing this are being discussed with stakeholders, for possible implementation during the 2003 data collection.

A Puerto Rico Community Survey with similar design and sampling rates, is planned starting in 2003, pending congressional funding.

A crucial part of the ACS message is that the ACS provides the *characteristics* of the population, not *counts*. The census will continue to provide a complete count of the population every ten years. In the intercensal years, the official number of people will continue to come from the intercensal demographic estimates program, as part of the Federal/State Cooperative Population Estimates (FSCPE) program. Information from the ACS will be used to improve these population estimates. Research has begun on these improvements and how to implement them as a "Program of Integrated Estimates."

To replace the long-form estimates, the ACS will produce annually updated 5-year average estimates for geographic areas down to the block group level. In 2008, for example, we plan data products covering the period 2003-2007. In 2009, the updated estimates will cover 2004-2008, and so forth. Each 5-year average may be thought of as replacing a hypothetical census long form in the middle year; for example, the 2003-2007 average would correspond to a "2005 long-form estimate." The 2008-2012 average is the one most closely corresponding to the 2010 time period. These updated 5-year averages are the most important ACS data product for small population groups because they will show the direction and level of trends, information never before available for smaller groups.

The ACS will also produce 3-year averages and 1-year average estimates. For larger areas and population groups of 20,000 or more people, there will be 3-year averages, and 1-year averages for areas and groups of 65,000 or more people. These averages will be regularly available and updated in the data products to show changes in characteristics over time.

For research purposes, we will make 3-year and 1-year averages available for smaller areas and population groups. We will discuss the details of the format (e.g., SAS files) with data users to determine what is generally the most useful. Research might include, for example, statistical analyses such as time series modeling or multiple regression analyses which pool information from a number of areas. These data will also be useful to help interpret multi-year averages, that is, to study in more detail the changes that took place within the 5-year period for which the averages are shown in the core products. The single year and 3-year averages for the smaller areas and groups will have high standard errors and are useful only for detecting large changes.

III. IMPROVED INTER-CENSAL ESTIMATES FOR RACE, ETHNIC, AND ANCESTRY GROUPS

The complete counts of race and Hispanic origin groups will be collected on the decennial census short form as always. The advantage the ACS provides is updated information about patterns of change in the size, social and economic characteristics, and geographic location of race, ethnic, and ancestry groups during the decade. This information will be incorporated into the intercensal population estimates program to improve their accuracy.

In years between censuses, the ACS offers clear improvements in the information available to estimate the number of people in the race and ethnic groups listed on the decennial census short form, including specific groups. This encompasses not only broad groups such as “Asian,” “Hispanic,” and “American Indian or Alaska Native,” but also specific groups such as “Korean,” “Jamaican,” “Puerto Rican,” “English,” “Cuban,” and specific American Indian tribes. For the decennial census, ancestry groups are collected on the sample (long) form only. The American Community Survey will update information about ancestry groups every year.

Information from the ACS, along with other advancements in the methodology for intercensal estimates, will improve the quality of sub-state estimates of the broad race/origin and ancestry groups. Without the ACS, the intercensal estimates program can now provide estimates only for the broad race groups or the total "Hispanic" category. There are no intercensal estimates for the detailed subgroups or specific American Indian tribes or for ancestry groups. Even for the broader groups, the intercensal estimates historically have not done well in reflecting changes in migration patterns below the state level.

To bring about improvements in the quality and detail of intercensal estimates for the smaller groups, we are developing an improved methodology that uses the ACS estimates of demographic characteristics. The ACS multi-year averages will provide information about changes between censuses in the characteristics and geographic areas in the detailed groups. As illustrated in subsequent sections, for small groups the ACS measures dramatic changes, the most important ones to measure. Appendix 1 includes some examples of dramatic change between 1990 and 2000, suggesting that large percentage changes occur frequently for smaller groups.

IV. RESPONSES TO QUESTIONS ABOUT THE ACS AND SMALL GROUPS

This discussion uses examples for 400 people to respond to questions that data users have asked about the ACS and small population groups. The examples could represent (1) the number of people in a particular population group in a particular area; or (2) the number of people in an area from a group who have a specific long-form characteristic, such as being employed in a particular industry, teenage mothers enrolled in school, or people who use a language other than English at home. This hypothetical example uses a relatively high, yet realistic, standard error for both the ACS and the long form. The relatively high standard error in the example would correspond to a characteristic that has the same value for all or most of the members of a household. Such characteristics tend to have higher-than-typical standard errors. As long form and ACS data become available for a wider range of characteristics, analyses like this one are being done using the actual standard errors for a variety of estimates, large and small.

Census stakeholders have asked us questions about the quality of estimates for small population groups from the American Community Survey. This section focuses on four basic questions and our responses to the concerns and how we are addressing the issues that data users have raised.

Question 1. *What is the impact of having a smaller ACS sample size in any single 5-year period than the long form has in the census year?*

Response: It is correct that a single 5-year average from the ACS is based on a smaller effective sample size[1] than the census long form. As such, the ACS estimates will have larger confidence intervals than long-form estimates. Since the long-form estimates

[1] The term effective sample size" refers to the number of distinct units in the sample and to the relative sample size that will result in a similar level of sampling error when compared with simple random sampling and unbiased estimation. Based on the ACS design, the "effective sample size" is about 64 percent.

already have large confidence intervals for small groups, this may make the data too noisy to be useful for some purposes but not for others. This is like the "blurry freeze frame" in our video metaphor – you may not be certain about what is occurring but you can get some hints, more information than having nothing at all. For a group of 400 people, the census long form would typically have a 90-percent confidence interval of roughly 280 - 520.[2] An ACS 5-year average would have a slightly larger interval, on the order of 240 - 560. In other words, a typical confidence interval for a hypothetical 2010 census long-form estimate of 400 would be ± 120. By comparison, a 2008-2012 ACS average estimate of 400 would have a 90-percent confidence interval of ± 160.

The basic premise of the ACS rolling sample is that this relatively moderate increase in the sampling error for one part of a decade is a reasonable tradeoff so as to profit from the ability to update the 5-year average every year and thereby gain a picture of the direction of change and relative differences among groups and areas. If the size and characteristics of the population change, such as from 400 to 480, the 5-year average gives a more accurate picture of current conditions than the out-of-date long form statistics (see Appendix 2). As shown in Appendix 1, small population groups change by much more than this. The updated ACS estimates would give a more accurate reflection of current conditions, compared with continuing to use the years-old previous census.

---

[3] This would be the confidence interval will be centered on 400, if the estimate is 400. The actual estimate would not be exactly equal to the population value because of sampling error. See Appendix 2-A. The length of the interval depends on what characteristic is being measured. See Appendix 2-B.

Another valuable aspect of the ACS is that it provides information about when during the decade changes take place. It helps us to move away from relying on national averages to imagine what is going on in regions of the country and among different population groups. For example, we can better identify geographic areas or groups that are in a recession when the nation is, on average, doing "well." Alternatively, we can identify areas of success when, on average, the nation is in a recession. This ability to go beyond national averages between censuses helps us assess the reasons for change and differences and how and whether the change is likely to continue. It could help decisionmakers to develop more proactive policies to prevent problems before they become serious. Because of the relatively small annual ACS sample size, this ability is limited to large changes, as discussed and illustrated in the next section.

Confidence intervals primarily reflect sampling error but also some aspects of nonsampling error. This discussion has not taken into account potential improvements in nonsampling error in the ACS due to experienced interviewers and follow up by telephone that result in more questions on the form being answered (that is, a higher "item" response rate). Such reductions in nonsampling errors compensate in part for the slightly larger confidence intervals compared with the decennial long form.

Question 2. *Are the multi-year averages more difficult to interpret than point-in-time estimates (a snapshot), especially when there are substantial changes in the population during the period of the estimate?*

Response. When there is little change in a population of less than 20,000 people, a single 5-year average is equivalent to a snapshot. When there is substantial change, the 5-year average is more like a blurry video for fast-moving objects that can be improved by the updated series of 5-year averages. The bottom line is that if the population is changing substantially, getting some information about the change is better than getting no information, as happens when data are collected only once in ten years.

A more detailed answer depends on the specific situation. In the next section, we provide examples of ways a population might change over time, and how to use such information from the ACS.

There are some cases where a single 5-year average only, with no updating, would not be as good as a decennial snapshot. It is the yearly updating that gives the ACS its advantage. In all the examples, the updated series of 5-year averages are preferable to statistics for only one year out of ten. In some situations, to make the most complete use of the ACS information, analysts would supplement the series of standard 5-year averages with information from the 3-year and 1-year averages in the "research files." Obviously, it would be ideal if we could collect the full long-form sample every year, but that is not an option because of the cost and public burden.

Question 3. *Won't there be an increase in the standard errors in areas where response by mail or telephone is relatively low because you use a subsample of 1 in 3* *nonresponse cases for follow up with personal visits from Field Representatives to collect the data?*

Response: Our evaluation studies show there are issues of both precision and differential bias among groups, and the Census Bureau is focusing on how we can best tackle these issues. It is important to reduce bias and bring the standard errors for all groups and areas in line with the objectives for response rates overall.

On average, about 60 percent of the population are represented by the ACS data collected by mail or CATI. For most of the remaining 40 percent of the population, the data are collected from a one-in-three subsample of nonrespondents. It is essential to maintain a low proportion of missing data for all areas and all population groups. Incomplete data increases the overall error of the estimates.

Aditionally, it is important to reduce statistical bias in the estimates for all groups. There is a potential for differential bias among groups if the survey systematically excludes people with characteristics that would be missed by the survey even if their address had been selected for sample and follow up.

From the evaluation studies we have learned that there is substantial variation in mail response rates by race and geography. Mail response rates in the testing phase have been lower for tracts with high proportions of African American or Hispanic populations. There is some evidence of substantially lower rates for tracts with high proportions of American Indian or Alaska Native population or Native Hawaiian and

Other Pacific Islander population. There is also evidence that households with limited English proficiency, including non-Hispanic households, have a lower-than-average mail return rate.

The Census Bureau proposes to address these issues through several techniques, including (a) using a subsample rate larger than 1 in 3 in areas with low mail response; and (b) making it easier for people with limited English proficiency to respond by mail or telephone. Nonsampling errors can be larger than sampling errors and so our research program monitors both. Because the Field Representatives are experienced, they have had good success in areas where it has traditionally been difficult to collect survey statistics during the short decennial census operations with temporary staff. This may partially compensate for subsampling the nonresponse cases for follow up to collect responses to the questionnaire.

Judging from 1990 census results, in areas and for population groups with lower-than-average mail response rates, the completeness of the long form data collection (that is, responses to all the questions) tends to be uneven as well. There is evidence from our early evaluations that the ACS has had more complete data collection for the units in its sample more consistently than is the case for the long form. While we continue to monitor and evaluate item nonresponse rates, we believe the higher completion rates in the ACS are because of its smaller, more experienced interviewing staff compared with the large number of temporary decennial census interviewers. Additionally, in the ACS, there is more opportunity over time to adjust our operations and methods and thereby improve data collection than is possible in the rushed environment of the decennial census data collection period.

Question 4: *How can a small monthly sample, such as that of the ACS, be representative of a small population group that is geographically dispersed?*

Response: In some months, it is possible that no one is selected from a particular small population group that is widely dispersed geographically. Even so, while the estimates for a single month may be very unpredictable, data averaged over 60 months provides reasonably stable estimates. Sampling statisticians use the laws of probability to select survey samples that are representative and have a specified margin of error. Intuitively, it is harder to visualize how the averages result in representative statistics when a population group is geographically scattered without any particular pattern and the sampling rate is relatively small. It is easier to visualize how a systematic sample, for example taking every sixth address, gives good representation for a population group that is clustered in a particular geographic area

These intuitive concerns about a "guarantee" of representativeness raise a legitimate issue. The laws of probability make "guarantees" only within a certain "margin of error" or "confidence interval." When the sample and population group are both small, the margin of error can be large, as a percentage of the survey estimate. The laws of probability do not ensure precise estimates from small samples. What the laws of probability do ensure is that statisticians can calculate how large the margin of error is likely to be due to sampling,[3] a topic covered in more detail in the next section.

Whether a survey's sample size is adequate depends on whether the confidence intervals for the survey estimates are small enough to allow data users to use the estimates for their purposes. A common way to think about the adequacy of confidence intervals is to consider how large a difference it would take in the survey's estimates to be "statistically significant." With census long-form estimates for two groups of about 400, each having a confidence interval of ± 120, the difference in the survey's estimates would not be statistically significant unless the two estimates were as different as about 315 for one group versus 485 for the other.

With the larger ACS confidence interval of ±160 for a 5-year average, the difference between averages of 315 and 485 would not be statistically significant. It would take a difference of 287 versus 513 to be statistically significant. This indicates the price paid because the proposed ACS has a smaller sample size in a single 5-year period than the long form has in the census year. And yet, it is not the *single* multi-year average that is the important comparison to make between the usefulness of the two data sets. It is the annually updated *series* of multi-year averages the allow data users to be better informed by understanding the level and direction of changes during the decade. Additionally, the reductions in nonsampling error in the ACS help to offset some of the "price" paid for slightly higher sampling errors.

[3] By contrast, it is hard to quantify how large the resulting error in the estimates is likely to be for nonsampling errors such as nonresponse, undercoverage, or misunderstanding of questions.

The ultimate question for users of statistics for small groups is whether the long form's slightly greater precision for comparing groups is of such practical importance that it is worth giving up the opportunity to learn about substantial changes in the size and characteristics of the small group over time. The premise of the ACS design is that, for small groups, the ability to learn about substantial changes over time is essential and worth a moderate loss of precision for any single point in time.

For example, consider the potential use of estimates of children under age 5 who speak a language other than English at home in helping school systems prepare and provide for appropriate educational opportunities in coming years. The *series* of ACS 5-year averages can monitor trends in the number of such children, and the 1- and 3-year averages can detect sudden large changes. By contrast, neither a single decennial estimate or a single 5-year average, whether 400±120 or 400±160, has the precision or timeliness to be much help in planning. The 2010 long-form statistics will be available in late 2012, in time for planning for the 2013-2014 school year and preferable to a single 2007-2011 ACS average that will also be available in mid-2012. Rather, it is the series of updated ACS averages that would alert school planners more quickly when there are large changes in the needs of children who will be entering the school system and thus better inform their strategic planning.

V. MORE DETAILED EXAMPLES AND DISCUSSION

Comparisons of standard errors and estimates between the long form and the American Community Survey

Below we will illustrate how the ACS standard errors for small populations change over time under different scenarios. For each figure, the population starts at 400 in the year 2010 but the graphic begins with the year 2012. First, we show the series of ACS 5-year averages and then the decennial long-form information that would be available to data users. In each figure:

- The “diamond” symbol represents the assumed actual value compared with the estimates from each data set;
- For the ACS graphics, the solid lines indicate the upper and lower bounds for the probable estimate averaged over the previous five years and then the updated survey estimate each year. For example, the bounds for the year 2018 show the range that has a 90-percent probability of containing the 2013-2017 average estimate for the ACS sample, given the population values indicated by the diamond symbols.

Figure 1 shows the series of ACS 5-year averaged estimates. Figure 2 shows the long-form estimates for a population that changes from 400 in the year 2010, to 1400 in the year 2020, and the corresponding probable actual values.

ACS 5-year Average (Figure 1)

Population with Strong Trend

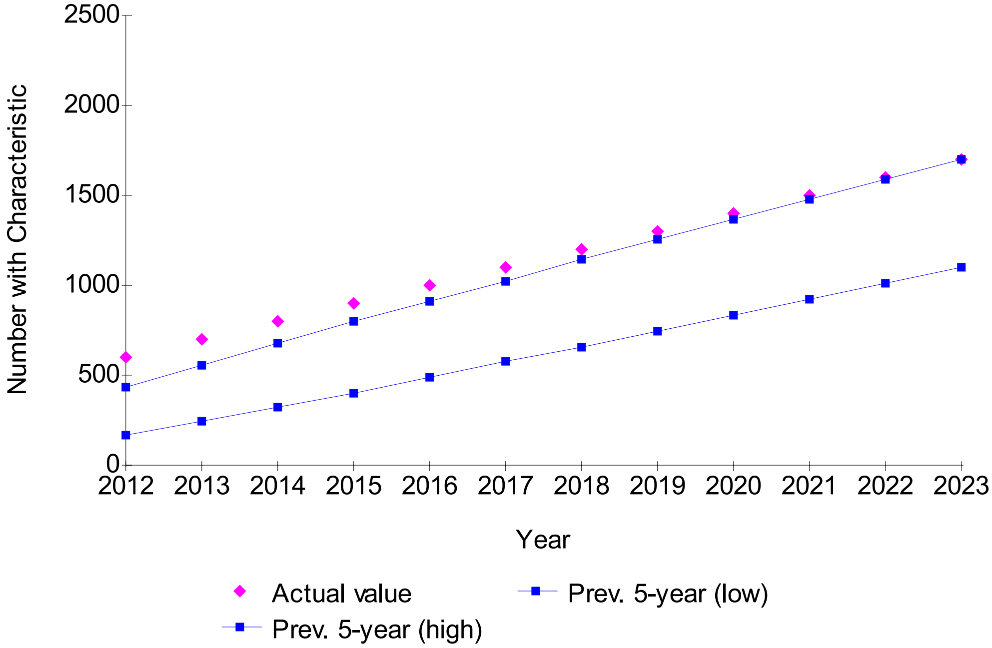

Decennial Long Form (Figure 2)

Population with Strong Trend

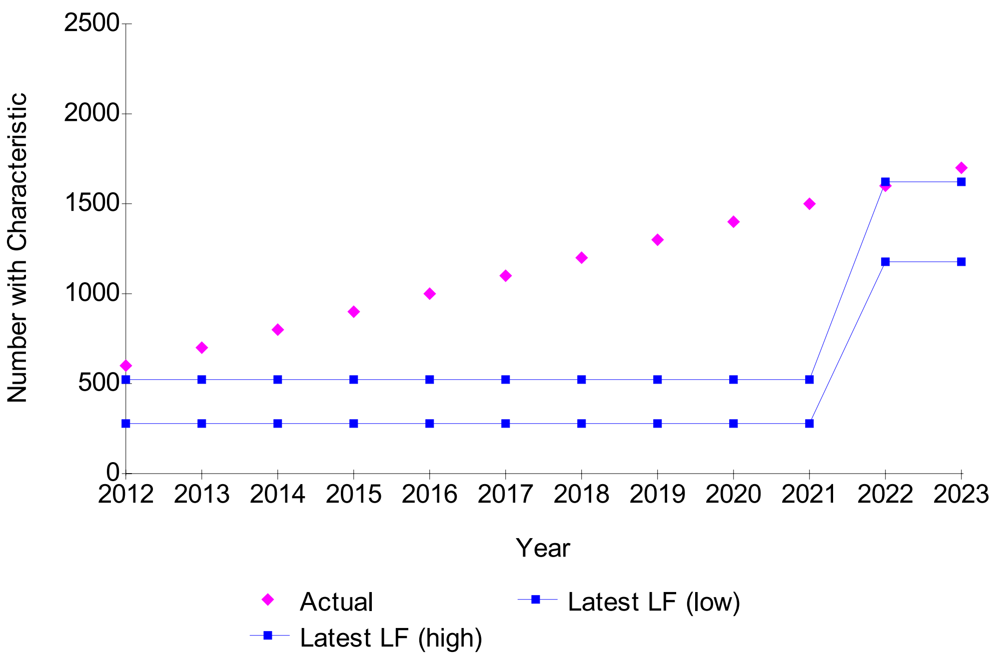

The increasing spread between the upper and lower bounds of the ACS estimates in Figure 1 occurs because the number of people with the characteristics is increasing. Larger estimates tend to have larger standard errors, although the standard error grows smaller as a percentage of the estimate.[4]

The 5-year averages in Figure 1 tend to lag slightly behind the actual population values, and the sampling errors are greater than those for the long form in Figure 2. Yet, the 5-year moving averages are obviously closer to the actual, current population value in most years than for the long form. Unlike the long form, the ACS 5-year averages reflect the direction of the actual trend.

Figures 3 and 4 show the same scenario using ACS 1- and 3-year averages. The 3-year averages in Figure 3 are a reasonable alternative to the 5-year averages for uses of the statistics where the smaller time lag would compensate for the higher sampling error. The single-year survey estimate in Figure 4 has a much larger range of probable error, and is not as useful unless the change is very large.

[4] For example, the standard error for the 2012 estimate is **. That is ** percent of the middle of the range of values (**). For the 2022 estimate, the standard error is ** and ** percent of the mid-range value of **.

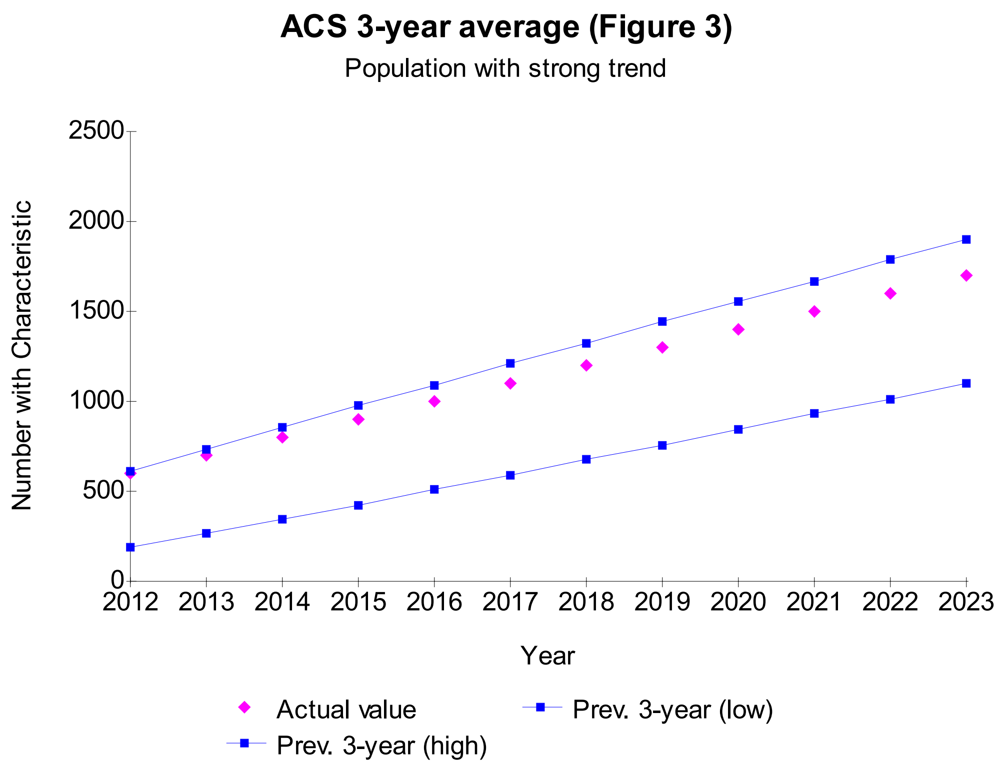
ACS 3-year average (Figure 3)
Population with strong trend
2500
2000
1500
1000
500
0
Number with Characteristic
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
Year
Actual value
Prev. 3-year (low)
Prev. 3-year (high)

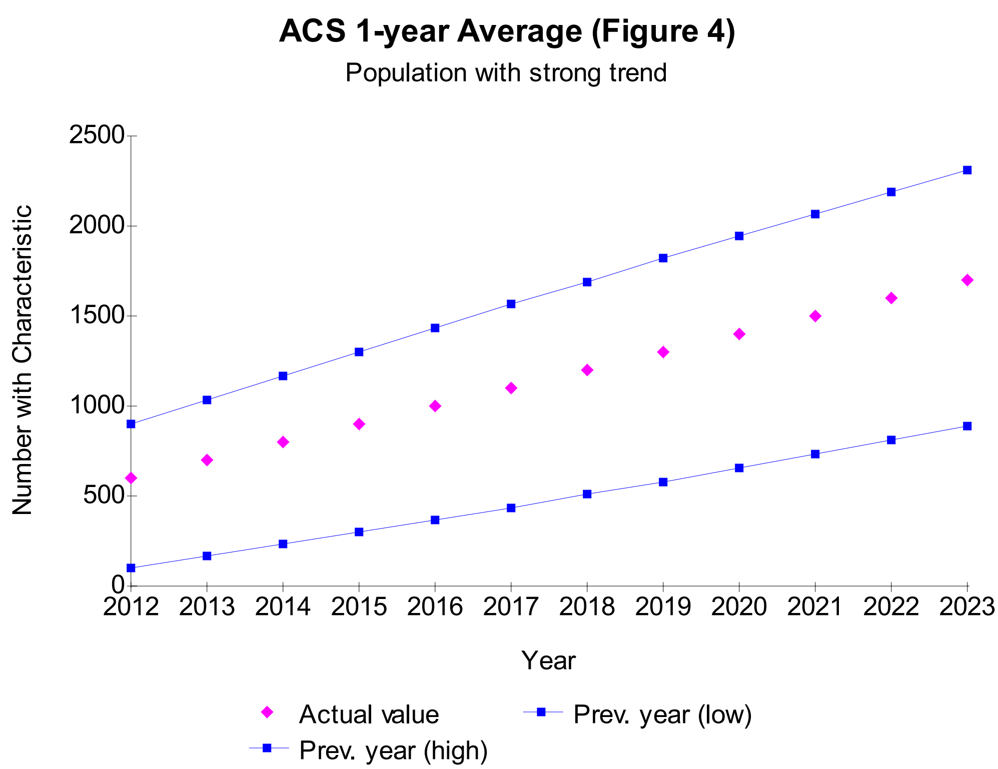
ACS 1-year Average (Figure 4)
Population with strong trend
2500
2000
1500
1000
500
0
Number with Characteristic
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
Year
Actual value
Prev. year (low)
Prev. year (high)

In Figures 5 through 8, the true values jump suddenly from 400 in 2010 to 1,400 at the end of 2014. The 5-year ACS averages in Figure 5 picks up the changes within a few years, much sooner than the decennial long form in Figure 6. The changes are fully reflected it in the 2015-2019 average. The 5-year averages give the impression, however, that there is a steady increase starting in 2015, rather than the sudden jump. This is not the best picture of the change, but still better than that provided by the long form's two points of information.

Figures 5 through 8 illustrate scenarios where the more detailed analysis using 3-year and 1-year averages is useful after seeing that the 5-year averages indicate an important change. In this extreme example, comparing each 1-year average to the previous year would give a good indication of the timing of the change. After learning from the 1-year numbers that there might be an unusual jump in 2015, the 3-year averages gives a better idea of the size of the jump without overly "smoothing" the change as the 5-year averages do. Having considered all three ACS charts, the changes (up or down), the analyst would know the direction and magnitude of the increase, and that it took place over a several years in the middle of the decade. The analyst might still be uncertain whether the change took place all in one year or over several years. None of this information would be available from two measurements taken ten years apart (Figure 6).

The practical implications for policy decision are obvious. The ACS allows informed decisions to be made in response to changing conditions. The decennial census documents two points of historical change after they have occurred over a decade.

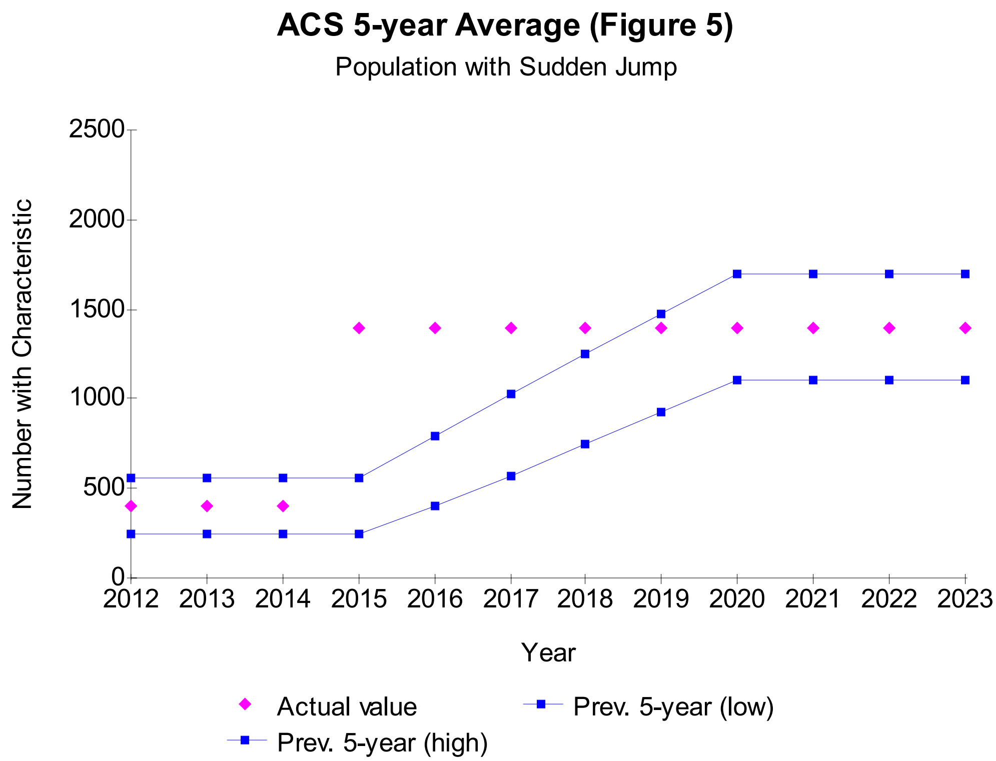
ACS 5-year Average (Figure 5)
Population with Sudden Jump
Number with Characteristic
2500
2000
1500
1000
500
0
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
Year
Actual value
Prev. 5-year (low)
Prev. 5-year (high)

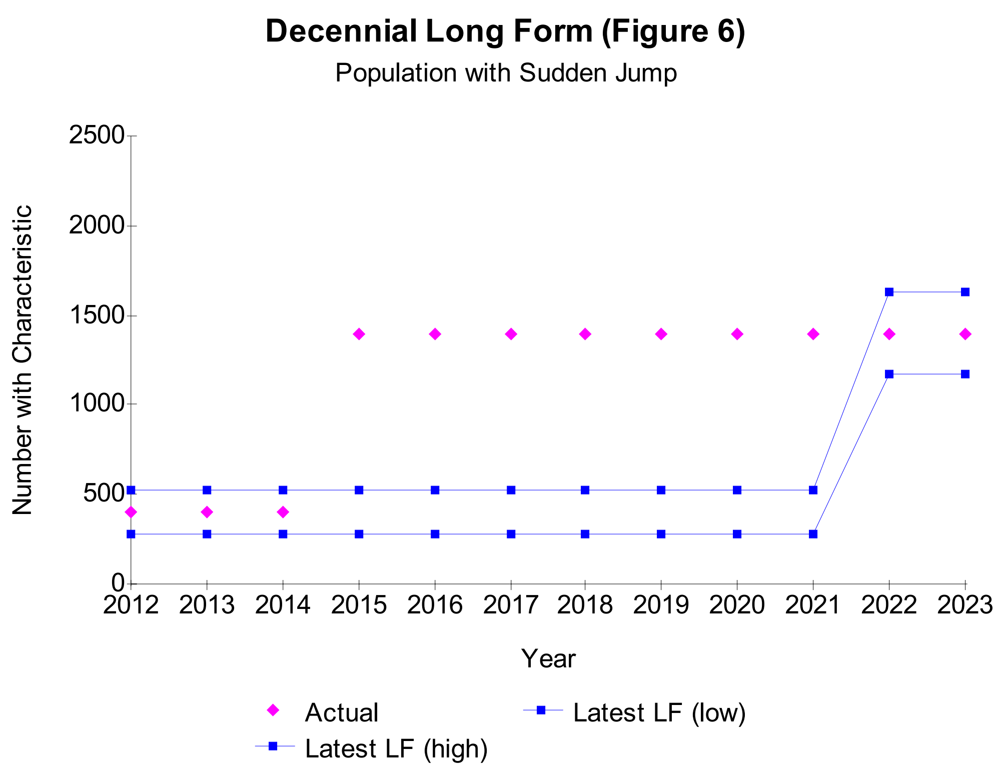
Decennial Long Form (Figure 6)
Population with Sudden Jump
Number with Characteristic
2500
2000
1500
1000
500
0
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
Year
Actual
Latest LF (low)
Latest LF (high)

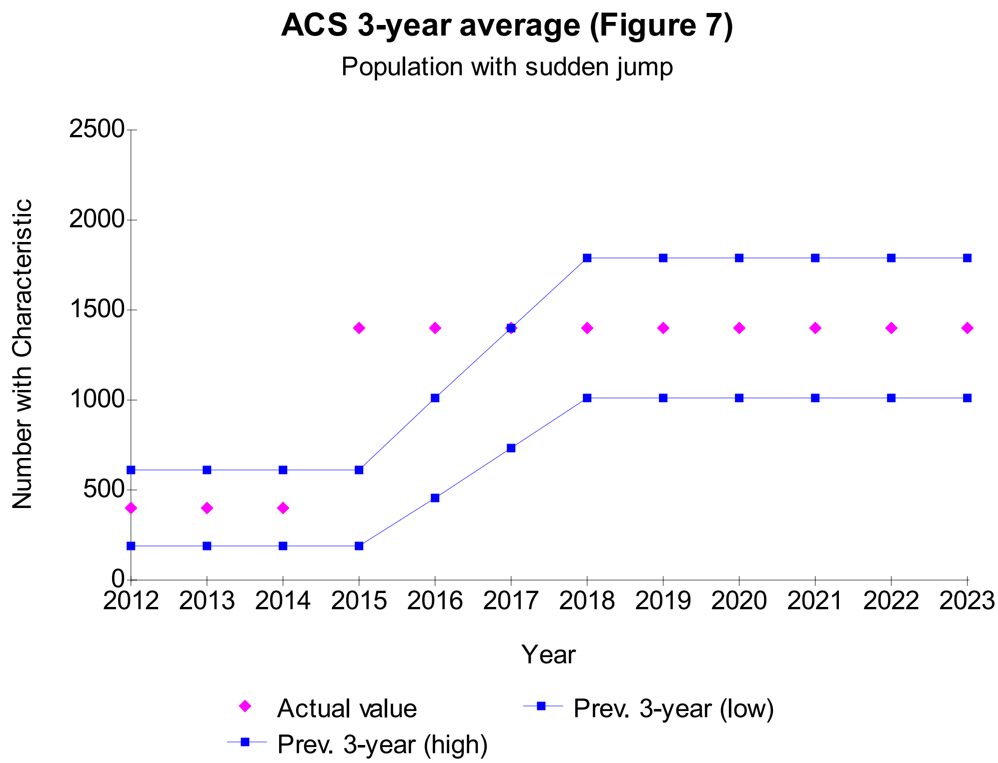
ACS 3-year average (Figure 7)
Population with sudden jump
2500
2000
1500
1000
500
0
Number with Characteristic
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
Year
Actual value
Prev. 3-year (low)
Prev. 3-year (high)

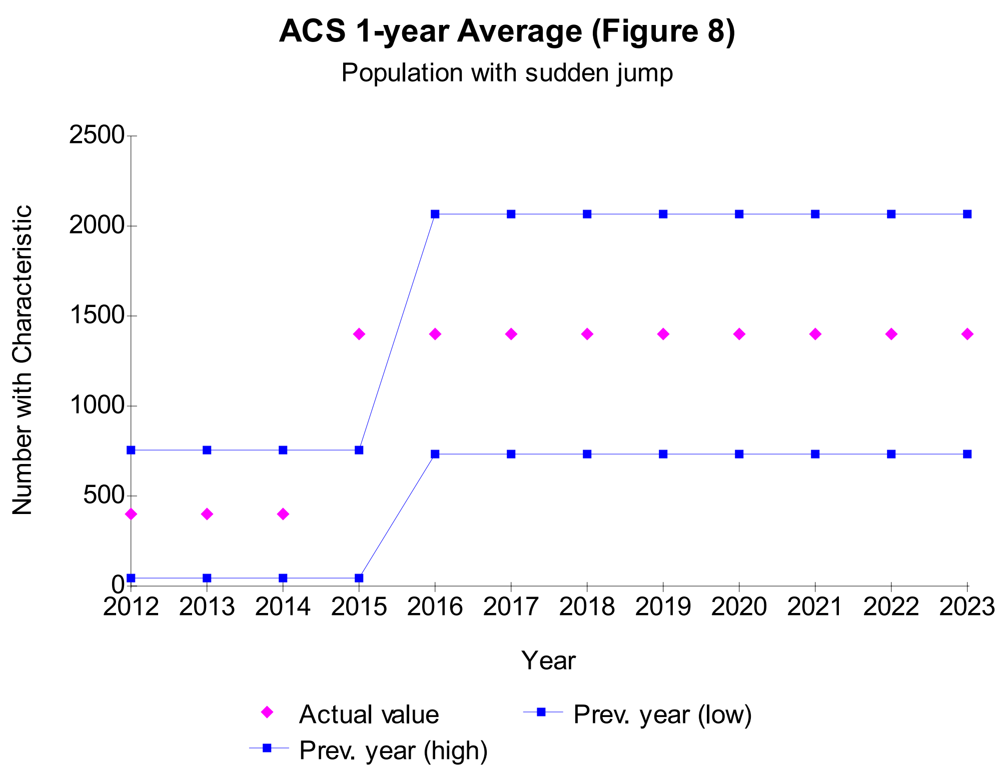
ACS 1-year Average (Figure 8)
Population with sudden jump
2500
2000
1500
1000
500
0
Number with Characteristic
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
Year
Actual value
Prev. year (low)
Prev. year (high)

Multi-year averages for changing populations.

If the population does not change meaningfully over a 5-year period, there is no issue about interpreting the 5-year average. For different patterns of change over time, as illustrated below, the average may relate in different ways to the single-year estimates. With the continuously collected ACS data, it is possible to get considerable information about the magnitude and direction of change over time. Because of the sampling error, however, it will not be possible to be sure of picking up a slight trend, or whether a strong trend is steady or somewhat irregular. The long form, of course, provides no trend information except for two points ten years apart.

The examples below address the question of how useful it would be to know only the information available from the averages, compared with knowing one individual value out of ten. To keep the examples simple, the tables below do not include the margins of error, as did the graphs Figure 1-8. Appendix 2 provides a discussion of some important statistical points for those who want a more detailed technical discussion.

In all the examples, averages that cannot be calculated from the data for the years shown in the tables are left blank to make the example easier to follow. These rules would be available from the ACS documentation once it has been fully implemented.

In most of the examples, the census year is the fifth year shown in the table, so data before and after the census are shown. In some examples, to illustrate what would have happened if the pattern of change had occurred one year earlier compared to the census, there is an additional row of numbers showing what would be measured by a census in the sixth year.

1. **Steady Trend.** If there is a steady increase or decrease in the size of the group being measured over a 5-year period, then the 5-year average corresponds to the value in the middle year of the average. For example, the second row shows that the average for years 1 through 5, which is available in year 6, is 440. This is equal to the actual size of the group in year 3 as, shown in the first row.

EXAMPLE 1: A STEADY TREND

Year (y)	1	2	3	4	5	6	7	8	9	10	11	12
Actual Size of group in year y	400	420	**440**	460	480	500	520	540	560	580	600	620
Average of previous 5 years	___	___	___	___	___	**440**	460	480	500	520	540	560
Previous census (year 5 census)	__	___	___	___	___	___	480	480	480	480	480	480

If the steady increase continues, the series of averages will give an accurate description of the trend, albeit with a three-year lag. The decennial snapshot, in the third row, misses the trend entirely, becomes steadily more out of date, and fails to measure the trend. If the trend of annual data is somewhat irregular, the moving average will tend to smooth out the irregularities, making the trend look steadier than it actually is. A smooth increase in the moving averages means that the actual population is generally trending upwards, but not necessarily as steadily as the averages suggest.

2. **Sudden jump or drop.** The 5-year averages will show an increase when there is a sudden jump, but they will smooth it out, masking the suddenness of the change. This is illustrated in Figures 5 through 8 above and in the second row of the following example. Two possible census years are shown in the example, a "year 5" census illustrating a jump which occurs right after the census, and a "year 6" census illustrating a jump which occurs right before the census.

EXAMPLE 2: SUDDEN JUMP

Year (y)	1	2	3	4	5	6	7	8	9	10	11	12
Actual Size of group in year y	400	400	400	400	400	600	600	600	600	600	600	600
Average of previous 5 years.	—	—	—	—	—	400	440	480	520	560	600	600
Previous census (year 5 census)	—	—	—	—	—	—	400	400	400	400	400	400
Previous census (year 6 census)	—	—	—	—	—	—	—	600	600	600	600	600

To detect the fact that the change is much more sudden than the 5-year averages indicate, it is necessary to look at the 3- and 1-year averages (Figures 7 and 8), as part of studying and interpreting the changes.

The uses of decennial census data, in a situation of sudden change depends very much on the year of the large change. If, as with the September 11 attacks in 2001, the census is taken shortly before an event of dramatic change, the census will provide a valuable profile of the area, but it will be 10 years before the impact of the event will be measured. This is illustrated in the third row of the table. If the census occurs after the change, as illustrated in the fourth row, it will instead provide a useful “after” profile. If the census occurs during a period of dramatic change, for example right after a natural disaster, the census may be disrupted by the event and the data may have limited value.

The ACS 5-year averages would give a baseline profile for the small group before the dramatic change in year 6. Eventually, it would give a 5-year average profile after the change (year 12). In that sense, it combines the information of a census before the change and a census after the change. As the example indicates, however, the picture given by the averages that cross the change year (years 7 through 11) requires careful interpretation.

The entire series of ACS moving averages have a clear advantage over a decennial snapshot in such situations. They provide before and after profiles of the area. The 1- and 3-year averages will give a useful earlier measure of the change if it is large.

3. **Irregular, seemingly patternless, change.** This seems at first to be the most difficult situation for interpreting an average. Actually, it is the most natural situation for using an average. Averages would often be used in such situations,

even if there were a census every year and there was no issue about sampling error. This is because the average, over a period of time, provides a more stable description of the area.

For this example, consider the populations in the group of interest in two areas (X and Y), each with considerable variation from year to year:

EXAMPLE 3A: IRREGULAR CHANGE POPULATION IN AREA X

Year (y)	1	2	3	4	5	6	7	8	9	10	11	12
Size of group in year y	159	263	226	367	117	253	79	298	234	64	159	162
Average of previous 5 years	—	—	—	—	—	226	245	208	223	196	186	167
Previous census (year 5 census)	—	—	—	—	—	—	117	117	117	117	117	117
Previous census (year 6 census)	—	—	—	—	—	—	—	253	253	253	253	253

EXAMPLE 3: IRREGULAR CHANGE POPULATION IN AREA Y

Year (y)	1	2	3	4	5	6	7	8	9	10	11	12
Size of group in year y	491	355	317	513	458	270	534	394	468	373	539	347
Average of previous 5 years	—	—	—	—	—	427	383	418	434	425	408	462
Previous census (year 5 census)	—	—	—	—	—	—	458	458	458	458	458	458
Previous census (year 6 census)	—	—	—	—	—	—	—	270	270	270	270	270

As the averages in the second row show, Area X's population averages around 200 while Area Y's population is generally higher at about 400. The 5-year averages in row 2 show this general relationship of the areas more clearly than the "unsmoothed" single year values in row 1.

A decennial snapshot can give a very misleading picture in such situations. In this example, if the snapshot is in year 5, it would give the impression of an unusually large difference between the areas (117 compared with 458 in row 3). If the census were in year 6, it would show the areas as being almost the same (253 compared with 270 in row 4).

4. **A single-year spike**. This is an extreme case where a decennial snapshot could be advantageous under very restricted circumstances, but in general the ACS information would be preferable, although far from perfect. In this example, members of a population move into an area for one year.

EXAMPLE 4: SINGLE-YEAR SPIKE

Year (y)	1	2	3	4	5	6	7	8	9	10	11	12
Size of group in year y	0	0	0	0	0	400	0	0	0	0	0	0
Average of previous 5 years	—	—	—	—	—	0	80	80	80	80	80	0
Previous census (year 5 census)	—	—	—	—	—	—	0	0	0	0	0	0
Previous census (year 6 census)	—	—	—	—	—	—	—	—	400	400	400	400

If the group happened to be in the area at census time, then there would be a full one-in-six sample to provide a description of the group with an error on the order of only ±120 as in line 4 of the example. This may not be a desirable outcome, unless there is independent information that the group was in the area for one year only (for example, in Alaska when a large group of single men came to build the oil pipeline and then most left when it was finished). In this case, the census data would do a poor job of describing the area over the next decade. If the group happened to be in the area outside of the census year, then the census would give no indication of the characteristics of the group, and no indication that it ever was in the area, as in line 3.

In this example, the ACS 5-year averages in line 2 would indicate that the group was in the area but would give a "blurred" picture of exactly when the group was there. The 5-year average would be something like 80 ± 72.[5] The 1-year estimates would give the most useful description of the situation. The ACS 1-year estimate would be based on a one-in-forty sample and would give a confidence interval something like 400 ± 360. The margin of error is the same percentage of the estimate for both the 1- and 5-year averages: 72 divided by 80 is the same as 360 divided by 400. This is because the 5-year average estimate in this example depends entirely on the data from the single year when the group was in the area. After looking at the 1-year series, the analyst would recognize the blurring effect of the 5-year averages, which gives a l-year estimate spread over five years.

---

[5] These intervals illustrate the margin of error and general magnitude of the estimate. An actual sample would probably give an estimate different than 80. See the discussion in Appendix 2, part A.

From the ACS, even though the 5-year averages give a confusing picture of the timing, the 1-year data would indicate that the group was in the area. The survey estimates of the characteristics of such a small group in an area for only one year will result in a large confidence interval. Generally, they are of little use unless the members of the group have very similar characteristics, such as being employed in a particular occupation.

In 9 years out of 10, a decennial snapshot would totally miss the group. If the census happens to be in the year the group is in the area, then it gives much better estimate of the group's characteristics for that single year than the ACS. In future years, however, continuing to use this snapshot would give an erroneous picture of the areas. So the decennial census would be advantageous only if the group is in the area during the census year and if there is independent information that this was a short-term event. Otherwise the ACS provides the more useful information.

Details on proposal for oversampling areas with low mail response

Specific proposals are given in the paper by Tersine, et. al. (See References at the end of this paper, 2002). There is increasing evidence of the completeness of data collected by the ACS field staff in areas with low mail response rates. This was seen in a study of the ACS statistics in the Bronx NY ACS test sites (Salvo and Lobo, 2002), as well as a U.S. Census Bureau study (Love, 2002). Studies for additional test sites will be available next year.

How does probability sampling work for small samples?

First, we provide a simplified explanation of probability sampling to illustrate basic principles. We then present a more exact theoretical model, and finally, describe how confidence intervals are calculated in practice.

A simplified explanation of probability sampling. Suppose the group of interest has 240 people. Each month the ACS selects a sample of about 1 in 480 people in the population, or about 1 in 40 for the year. Our expectation is that the sample will have about one person from the group in every two months of sample, or an expected "one-half person each month" (240/480).

With 60 months of sample, this is statistically analogous to tossing 60 coins, each with a fifty-fifty chance of heads coming up on each toss of the coin. That is, we theoretically expect heads to come up 30 times and tails to come up 30 times. The total number of people from the group who fall into sample in the 60 months corresponds to the total number of heads up from the 60 coin tosses.

To move from the number of people in the group who are selected to be in the sample, to publishing the estimated number of people in the group who are in the population, the survey data are weighted. Weighting is a statistical procedure that adjusts the sample for the groups or areas with different chances of selection. That way, the survey estimates represent the smaller demographic groups and geographic areas that may have had an unequal chance of selection. (See References, Dahl). The details for what factors are used to weight each year of the American Community Survey are on the American Community Survey website under "Accuracy of the Data" at: www.census.gov/acs/www/AdvMeth/Accuracy/Accuracy1.htm

In its simplest form, "weighting" is multiplying the people (or households) in the sample by their probability of selection (the "basic" weight) so that the weighted survey estimates equal the Census Bureau's population and housing units counts (or estimates between censuses) for each geographic area down to the county level. Because every survey has nonsampling errors, we also adjust for errors we know about (e.g., noninterview adjustment). Since the ACS samples about one person in 40 every year, each sample person in any given year is counted as 40 people in the population for that year's estimate. To produce the 5-year average, the five annual estimates are added together and divided by five. This means that the 5-year average would be eight times the number of people in sample in the 5-year period (counted as 40 in the annual estimate and then divided by five to get the 5-year average).

This means that if exactly 30 of the 60 "coin tosses" were "heads up," then the 5-year average estimate would be 8x30=240, the "expected" result. If a large number of different samples were selected, this would be the average result from all the sample estimates combined.

Moving away from theory to reality, with 60 coin tosses, there might not be exactly 30 heads up. The range of results that are likely to occur, however, is predictable. There is a 90-percent probability that out of 60 coin tosses, the number of heads will be at least 23 but no more than 37 (that is, as we express it, the “confidence interval” is 23 – 37).

Likewise, the exact results from 60 months of ACS sample may not be the “expected” number of 240, but the likely range of the result for each statistic can be predicted. Multiplying by 8 gives the estimated number of people with the specified characteristic in the range of 184 to 296. That is, there is a 10-percent chance that the estimate will be outside this range. When that happens, it is rarely far outside the range. In tossing 60 coins, there is a 99-percent probability that the weighted estimate will be in the range of 160 to 320, and a 99.9-percent probability that it will be in the range of 136 to 344.

This predictability of the range of results that will be obtained from a large number of random outcomes is the basis for the science of sample surveys. It is the reason that sample units are selected using random numbers rather than someone’s judgment of what a good sample would be. The same principle is used in applications as varied as actuarial life tables, random clinical trials, and predicting payouts at casinos. Anyone who would like to see how random sampling works is encouraged to actually perform the experiment of shaking 60 coins in a cup, pouring them out onto a table, and counting the number of heads and tails that come up. Repeat this five to ten times to see for yourself how the results vary but stay within a predictable range.

A more realistic model. The number of people in a small population group who are part of the ACS sample is, of course, not produced by tossing a coin to get either one person or zero people in the sample each month. The 400 people in the group of interest live in households or group quarters. Some living quarters have one person from the group, some have two or more. Some households will not respond by mail or telephone. For nonresponding households, there is a two-thirds chance that it will be dropped from the sample and a one-third chance that it will be selected for a follow-up visit to collect responses for each household member. The results are then weighted to count in the estimates as three households with the same characteristics as the sample household from which responses were collected.

A more appropriate model would be what probability theorists sometimes call an "urn model." Imagine filling an urn with as many beads as there are addresses. Most of the beads are labeled "zero," but some are labeled with one, two, or some other number, with the sum of the numbers equaling 400. Some of the beads are labeled "N" for "nonresponse by mail or telephone." A sample of 1 in of 480 beads is selected each month. For each one of the sample beads which is labeled "N," a random number is chosen to give a one-third chance of keeping the bead in the sample, and multiplying its number by three. The remaining beads have a two-thirds chance of being removed from the sample.

Mathematicians have analyzed theoretical experiments like this, and also confirmed their analysis with small-scale experiments with actual beads and boxes, and larger-scale experiments with computer-generated random numbers. Such experiments show the same sort of predictability of results as coin tosses. The exact numerical results giving the likely range of values are more complicated to work out than for coin tosses, but for reasonably large samples, the same mathematical approximations can be used to work out the likely range of values for the estimates from a sample of 480 “beads” from any such “urn.”

The same rules of arithmetic that apply to beads with numbers on them also apply to households with a number of people in each one. We assign an identification label to each address, and use a random process to select which labels are in sample. The numerical results of the sample of households from a list of addresses are nearly the same as if beads were being drawn at random from urns.

Modern methods for estimating standard errors. Up until the 1960s, survey confidence intervals were estimated by learning enough about the population to deduce the standard error theoretically from mathematical models, such as the urn model described above. For each characteristic, it was necessary to make assumptions about the pattern of “clustering” of the characteristic within households. Highly clustered characteristics such as ancestry, where all members of a household often have the same characteristic, tend to have higher standard errors than less clustered characteristics such as people with disabilities.

Since the 1970s, the more usual procedure has been to use “replication methods.” These methods work by splitting the sample into pieces, so that each piece is a microcosm or “replicate” of the full sample. Then the variation among the pieces is measured, and the standard errors are mathematically deduced from the measured variation among the pieces.

Mathematical theory, and tests with simulated populations, show that these methods give numerical results similar to the more traditional "urn model" approaches for deriving the standard errors. The replication methods can more readily reflect the authentic clustering patterns found in the real populations, because they are calculated from the actual sample data, not a theoretical model. The ACS and the census long form both use replication methods to calculate standard errors. As a result, the standard errors reflect the actual clustering found for various population groups.

In summary, probability sampling works because of the predictability of random events in the aggregate. No single month of the ACS sample is very predictable, but 60 months is enough of an aggregate for the results to be predictable within a calculatable margin of error. The calculation of error is based on statistical methods that have been developed to deal with random events that are more complicated than, but basically similar to tossing coins and counting the number of times heads come up.

Appendix 1

Examples of Large Growth For Small Population Groups in the ACS Comparison Counties

Population Group	County	1990 Estimate	2000 Estimate	Sources (1990/2000)
Asian Indian	Pima, AZ	1,041	2,105	STF-1/SF-1
Chinese	Ft Bend, TX	4,072	10,500	STF-1/SF-1
Korean	Lake, IL	1,923	4,089	STF-1/SF-1
Vietnamese	Douglas, NE	529	1,122	STF-1/SF-1
Black or African American	Schuylkill, PA	842	3,147	STF-1/SF-1
American Indian or Alaska Native	Bronx NY	6,069	11,371	STF-1/SF-1
American Indian or Alaska Native	Lake, IL	1,198	1,801	STF-1/SF-1
Native Hawaiian and Other Pacific Islander	Bronx NY	541	1,383	STF-1/SF-1
Other Micronesian	Multnomah, OR	181	505	STF-1/SF-1
Dominican	Broward, FL	3,489	8,869	STF-3/ACS
Salvadoran	Douglas, NE	52	414	STF-3/ACS
Arab	Broward, FL	5,174	9,461	STF-3/ACS
Ukrainian	Multnomah, OR	1,524	5,469	STF-3/ACS

NOTES: The 2000 estimates for race or Hispanic origin are for those marking one race or one origin. The ancestry estimates are the first ancestry reported on the form. Census counts have been used when available. For the "Native Hawaiian and Other Pacific Islander" count for 1990, the detailed race tables were used. For the ACS estimates, the lower bound of the confidence interval for the 2000 data is shown. This is the most conservative estimate, and the actual growth is likely to have been larger than that shown in the tables.

Appendix 2
Technical Discussion of Statistical Issues

A. Population values and estimates. The discussions in Sections IV and V gloss over an important distinction made in theoretical statistics. In Figures 1-8, the graphs illustrate the range of likely estimated values, *assuming that the population has 400 people in the group of interest*. In Section IV, the discussion considered what margin of error would be associated with a survey estimate, *assuming that the estimated number of people was 400*. Since the goal of the discussion was to give only a general idea of how the ACS and long form standard errors compare, and not a statistics lesson, the subtle distinction was not emphasized.

The distinction will be illustrated with an example. Suppose that the population of the group is a constant 400 people for the years 2008-2012, so that the actual 5-year average population is 400 and also, the 2010 count would be 400. The estimate from a long-form sample could range, with a 90-percent probability, from 280 to 520. The corresponding ACS range is 240 to 560. This range is what the graphs in Section V illustrate.

The actual confidence interval would be centered around the estimated value, not around the population value of 400. If the long form gave an estimate at the upper end of its range, the confidence interval would be 520±137 (383 to 657). If the ACS estimate were at the upper end of its range, it would be 560±189 (371 to 749). Alternatively, the ACS or the long form could produce estimates at the low end of the range. Exactly where the estimated values fall depends on which sample of addresses happens to be selected. The size of the margin of error depends on the size of the estimated value as well as the sample size.

The example in the text simplistically assumed that the estimated value was 400. The midpoint of the range of estimates is 400, so in that sense 400 is a "typical" value for an estimate, and 400±120 or 400±160 is used to illustrate a typical confidence interval. However, the actual situation is more complicated as indicated in the previous paragraph.

B. Total Error. The total survey error is usually partitioned into "variance" and "bias." The variance refers mainly to the sampling error, and is expressed by the margin of error or the confidence interval. "Bias" is defined as the difference between the "true" value and the average of all possible samples.

In the example for question 1 (see second paragraph in Section IV), the long form estimate is biased because it was based on a sample selected from a population where the value was 400. That is, 400 is the average of all possible samples but the "true" value had grown to 480. In this example, the ACS has no bias as far as estimating the updated value of 480, but it has a larger standard error.

The usual measure of total error is the Mean Squared Error (MSE) given by the formula

$$\text{MSE} = (\text{Standard error})^2 + (\text{Bias})^2$$

At some point, as the population grows beyond the initial value of 400, and the bias in the outdated estimate grows, the value of $(\text{Bias})^2$ outweighs the value of $(\text{Standard Error})^2$. As a result, the Mean Squared Error of the long-form estimate is larger than that of the ACS. In this example, 480 is just beyond the break-even point.

Of course, if the "truth" is defined as the most recent population value rather than the previous 5-year average, then the ACS 5-year average also has a bias because of population growth during the 5 years. Defining the "truth" this way increases the MSE of the ACS. It tends to increase the MSE of the outdated long form even more, because the bias term is squared. As an example of the last point, consider the steady trend in Example 1 of Section V. By year 10, five years after the census, the most recent 5-year average would be 520. Compared with this, the census has a "bias" of –40 = 480-520, or a squared bias of 1600. If the "truth" is assumed to be 520, the 5-year average has zero bias. If the "truth" is taken to be the actual value in year 10, which is 600, than the 5-year average has a bias of –80=520-600 or a squared bias of 6,400. The census has a bias of –120=480-600, giving a squared bias of 14,400. The squared bias for the census therefore, exceeds the squared bias of the 5-year average by 8,000 rather than 1,600.

C. "Margin of Error" and "Confidence Interval." The term "90-percent confidence interval" is defined in the statistical literature as an interval from sample data, calculated in such a way that there is a 90-percent probability that the interval will contain the population value. A longer interval is the "95-percent confidence interval." For that confidence interval, there is a 95-percent probability of containing the actual population value. The interval can be expressed[1] either in the form 400±160 or as 240 to 560.

The term "margin of error" has several meanings. It is sometimes used as a general term, synonymous with "the plus or minus" amount in any confidence interval. At other times, its meaning is restricted to political polling. The term is also used as a general reference to uncertainty in the estimate, with no specific numerical measure in mind.

In political polling, the margin of error for a poll is the margin of error for a 95-percent confidence interval, assuming that each candidate has 50 percent of the vote. The pollsters give one margin of error for the entire poll, rather than a different margin of error for each estimate. For a typical poll, a candidate polling at 50 percent might have a confidence interval of 50±3, while one polling at 10 percent would have an interval of 10±2. In this case, the poll would be said to have a margin of error of "3 points, plus or minus" because that is the largest margin of error.[7] In the more general use of the term, the first estimate would have a margin of error of ±3, and the second estimate, a margin of error of ±2.

---

[7] Estimates greater than 50 percent have a smaller margin of error that the 50 percent estimates. For example, an estimate of 90 percent would have a confidence interval of 90±2.

Bibliography

Alexander, Charles H.. Jr., "American Community Survey Data for Economic Analysis," paper presented at the Census Advisory Committee of the American Economic Association, October 18, 2001. Available at: http://www.census.gov/acs/www/AdvMeth/Papers/Papers1.htm

Scott Boggess & Jonathan Chong, "Skills and Abilities of the Working-Age Population, Rockland County, NY: 1996-1999," Population Division, Education and Social Stratification Branch. April 2001, Available at: http://www.census.gov/acs/www/AdvMeth/Papers/Papers1.htm

Dahl, Scott A., "Weighting the 1996 and 1997 American Community Surveys," March 1998. Available at: http://www.census.gov/acs/www/AdvMeth/Papers/ACS/Paper12.htm

Kalton, Graham, Daniel B. Levine, Joseph Waksberg, John Helmick (Westat), "The American Community Survey and the Quality of Rural Data," June 29, 1998. Available at: http://www.census.gov/acs/www/AdvMeth/Papers/Papers1.htm

Love, Susan P., "Difficulties of Comparing Quality Measures Between the Census and the American Community Survey," U.S. Census Bureau, discussion from Power Point presentation at the 2002 Joint Statistical meetings.

Salvo, Joseph J and Arun Peter Lobo, "The American Community Survey: Quality of Response by Mode of Data Collection in the Bronx Test Site," Population Division, New York City Department of City Planning. Paper presented at the 2002 Joint Statistical meetings.

Salvo, Joseph J and Arun Peter Lobo, "The American Community Survey: Nonresponse Follow-Up in the Rockland County Test Site," Presented to the Annual Meeting of the American Statistical Association, Anaheim, California, August 1997. Available at: http://www.census.gov/acs/www/AdvMeth/Papers/ACS/Paper10.htm

Taeuber, Cynthia and Cheryl Chambers, "Fresh Data Every Year: The American Community Survey Can Help Improve the Measure of Poverty," Poverty Research News, Summer 1998, Available at: http://www.census.gov/acs/www/AdvMeth/Papers/Papers1.htm

Tersine, Anthony, "Item Nonresponse: 1996 American Community Survey," March 1998, Available at: http://www.census.gov/acs/www/AdvMeth/Papers/ACS/Paper13.htm

Tersine, Anthony, et.al., "Alternative Oversampling Options for Low Mail Response Areas in the American Community Survey," presented at 2002 Joint Statistical Meetings. Available at: http://www.census.gov/acs/www/AdvMeth/Papers/Papers1.htm